

ETHERNET BASICS

B&B ELECTRONICS

Regular, Fast, and Ultrafast Ethernet

10BASE5—Thick Ethernet or Thicknet—is the original Institute of Electronic and Electrical Engineers (IEEE) 802.3 Ethernet.

10BASE2—Thin Ethernet, Thinnet, or Cheapernet—resembles 10BASE5, was introduced to reduce the cost and complexity of installation, and became very popular for replacing Thick Ethernet as an office-cabling solution.

10BASE-T, a completely new physical layer, is IEEE 802.3i and uses two pairs of unshielded twisted-pair (UTP) telephone-type cable: one to transmit; the other to receive.

10BASE-F refers to three different fiber-optic specifications:

- **10BASE-FL** (Fiber Link) replaces the 1987 Fiber Optic Inter-Repeater Link (FOIRL) specification and is backward compatible with existing FOIRL devices. It is the most popular 10-Mbps fiber standard.
- **10BASE-FP** and **10BASE-FB** are dead. P means passive; B means backbone.

Fast Ethernet 100BASE-T is 10BASE-T with the original Ethernet Media Access Controller (MAC) at 10 times the speed. It allows three physical-layer implementations, all part of IEEE 802.3u: *100BASE-TX*, which has two pairs of Category 5 UTP or Type 1 STP cabling and is most popular for *horizontal* connections; *100BASE-FX*, which has two strands of multimode fiber and is most popular for *vertical* or backbone connections; *100BASE-T4*, which has four pairs of Category 3 or better cabling and is not common.

Gigabit or 1000-Mb Ethernet is the 1998 IEEE 802.3z standard that includes the Gigabit Ethernet MAC and three physical layers. Gigabit uses 8B/10B encoding and encompasses three physical standards: *1000BASE-SX Fiber* (horizontal fiber), *1000BASE-LX Fiber* (vertical or campus backbone), *1000BASE-CX Copper* (Copper-Twinax cabling), and *1000BASE-T*.

What Means What?

The IEEE naming convention for Ethernet is this:

- The first number (10, 100, 1000) indicates the transmission speed in megabits per second.
- The second term indicates transmission type: BASE = baseband; BROAD = broadband.
- The last number indicates segment length. A 5 means a 500-meter (500-m) segment length from original Thicknet.

In newer IEEE standards, letters replace numbers. For example, in 10BASE-T, the T means unshielded twisted-pair cables; in 100BASE-T4, the T4 indicates four such pairs.

7-Layer Network Concept

[Adapted from *Sensors Magazine*, July 2001 ©Advanstar]

For many years the ISO/OSI model has described the layers of information in a network, particularly the low-level transport mechanisms. From top to bottom, these are the layers and how these layers relate to your product design.


Layer	Name	Function
7	Application	Meaning of data
6	Presentation	Building blocks of data and encryption
5	Session	Opening and closing of specific communication paths
4	Transport	Error checking
3	Network	Determination of data paths within the network
2	Data Link	Data transmission, source, destination, and checksum
1	Physical	Voltage levels, signal connections, wire, or fiber

Most networks do not use all layers. For example, Ethernet and RS-232 are physical layers: layer 1 only for RS-232; layers 1 and 2 for Ethernet. TCP/IP is a protocol, not a network, and uses layers 3 and 4 regardless of whether layers 1 and 2 are a telephone line, wireless connection, or 10BASE-T Ethernet cable.

Ethernet Network Design Rule

The “5-4-3-2” rule states that the maximum transmission path is composed of five segments linked by four repeaters and, at most, the segments can be made of three coaxial segments with station nodes and two link [10BASE-FL] segments with no nodes between. Exceeding these rules means that some nodes will be unable to communicate with some other nodes.

This is a popular simplification of the IEEE 802.3 rule regarding maximum transmission length between two nodes:



(Note: these rules only apply when using hubs, not switches or routers.)

Building Blocks

An Ethernet network can have up to 1,024 nodes, hundreds of cables, and infinite combinations of hubs, switches, bridges, routers, network interface cards, and servers.

Hubs, the simplest method of redistributing data, are “dumb,” not interpreting or sorting messages that pass through them. A hub can be as simple as an electrical buffer with simple noise filtering. It isolates the impedances of multiple spokes in a star topology. Some hubs also have limited store-and-forward capability. They indiscriminately transmit data to all other devices, which are still on the same collision domain, connected to the hubs. They are not assigned MAC addresses or IP addresses.

Types: *workgroup*, usually stand-alone units with four to eight ports; *stackable*, with many more ports that can be linked to form a “super hub;” *segmented* which allow available ports to be divided among multiple groups and collision domains; *two-speed* which allow multiple baud rates to operate on the same hub, auto-detecting the data rate at each port and linking ports together with a speed-matching bridge; *managed* which have modest levels of intelligence and can be controlled remotely via a configuration port; and *repeaters*, essentially two-port hubs, which clean up the signal and boost it over distances.

Stacking or crossover cables allow multiple hubs to be connected in a daisy-chain topology. You can't use standard cables to link two hubs or two NIC cards together because they link transmit pins to transmit pins instead of transmit pins to receive pins. Note: Some hubs have *crossover ports*, which allow standard cables to be used.

Bridges, which operate at layer 2 of the OSI model, effectively extend the reach of each segment and allow traffic to selectively pass between two network segments. Bridges make forwarding decisions based on MAC addresses.

Intelligent bridges learn over time what devices are connected on each side and figure out which messages to forward and which ones to block. Such a bridge will automatically adapt to changes made to the networks over time.

Don't form loops on networks with multiple connecting bridges. The IEEE 802.1 spanning-tree algorithm removes loops: One bridge in a loop becomes the root and all other bridges send frames to it. Note: Bridges are not assigned MAC or IP addresses.

Routers, which maintain tables of IP addresses on each segment, learn the most direct paths for sending data packets to their destination. Routers are protocol-dependent because they operate at layer 3. Types include: 2-port; multi-port; access, which use modems; bridging, also called brouters, which change from router to bridge when they receive a packet they don't understand and send the message.

Terminal servers connect multiple serial devices—RS-232, -422, -485—to Ethernet. Types include *thin*, which link a single device to Ethernet and allow COM ports on the other side of the network to appear as though they are local to your PC; and *network time servers*, which use a global positioning system (GPS) to provide accurate local time for synchronization of devices and time stamping of events.

Gateways convert messages from one protocol to another, such as Modbus on RS-232 to Ethernet Modbus/TCP, or DeviceNet to EtherNet/IP. In most cases, the physical layers, protocols, and speeds are different. Gateways normally must be configured to work properly and are temporary rather than permanent solutions. Interface cards (NICs) link the PC to Ethernet via the PCI, ISA, PCMCIA, PC/104, or other buses. NICs handle layers 1 and 2, while the host processor in the PC handles everything else.

Software handles all other network layers, including TCP/IP, which comes built-in to nearly all PC operating systems, including Windows, Linux, DOS, UNIX, VxWorks, etc. Software applications for control or operator interfaces have drivers that pass application data, including higher-layer protocols like Modbus/TCP and EtherNet/IP to TCP/IP.

Note: There are no defined standards for *driver performance*—and there can be significant performance issues and delays with respect to driver and application performance. Many drivers are not written to serve the needs of deterministic applications. Response time can vary considerably on the basis of CPU speed, memory, how well the drivers are written, what other applications are running on the PC, etc. In most cases, the speed of industrial Ethernet networks will be limited by software and drivers, not by Ethernet itself.

Selecting the Right Cable

Select the right cable—and ensure that all components and interconnects on the network must also be equal to the cable's quality level. Grounding coaxial cable is generally good: It dissipates static electricity and makes your network safer. Use fiber-optic cables to link buildings, not copper.

Shielded twisted pair (STP) cable is naturally more noise immune and is preferable to unshielded twisted pair or UTP in noisy situations. STP should have at least 40 dB CMRR and less than 0.1-pF capacitance unbalance per foot. Ground STP cable, making sure the ground is connected only at one end. CAT5 STP patch panels normally provide a grounding strip or bar.

Hubs and switches don't provide grounding—use cables.

It's wise to be pessimistic about cables' ability to reject noise from 220 VAC and 440 VAC power lines and noisy power supplies of a factory. Capacitance imbalance greater than 70 pF per 100 m can introduce harmonic distortion, resulting in bit errors.

The cost of cable is quite small compared to total equipment cost, so if you're looking to save money, this is not a place to do it. Choose a well-designed cable to minimize bit-error rate after installation—and that will give faster throughput and fewer glitches.

Fiber-optic cable—certainly more expensive but bypassing the electrical issues such as noise and harmonic distortion, grounding, etc., especially in high-speed networks—is a very attractive choice.

If one section of a network is exposed to excess amounts of electrical noise, it's best to isolate that section with switches—but if you must use a hub in a noisy environment, use one with some level of intelligence instead of a buffer.

If your equipment is subject to washdown or exposure to corrosive chemicals, select cables with insulation rated to withstand exposure to those chemicals, such as PUR (polyurethane).

Choosing the Right, Best Network

Ethernet must work with other network technologies—and, for some applications, other networks will deliver more cost-effective performance. Ask these:

- What is the distance requirement?
- What physical cabling arrangement makes sense for this application?
- What is the speed (i.e., response time) requirement for the most time-critical devices? Do all devices require that speed or should some have a higher priority than others?
- Does the network allow you to prioritize messages?
- Do the devices you want to use support the same network standard? Are there open versus closed architecture considerations?
- If you are developing a network-capable product, what is the hardware bill of materials and cost of software development for that network?
- How much electrical noise is present in the application and how susceptible is the cabling to such interference?
- What is the maximum required packet size for the data that will be sent? If the data can be fragmented over several packets, how fast does a completed message have to arrive?
- What type(s) of device relationships are desired: master/slave, peer-peer, or broadcast?
- Does the network need to distribute electrical power? If so, how much current?
- What kind of fault tolerance needs to be built into the network architecture?
- What is the total estimated installed cost?

This paper includes excerpts from *Industrial Ethernet—A Pocket Guide. How to Plan, Install, and Maintain TCP/IP Ethernet Networks: The Basic Reference Guide for Automation and Process Control Engineers*. By Perry S. Marshall. ISA. 2002. Research Triangle Park, N.C. ISBN 1-55617-788-7.